An Efficient Summary Graph Driven Method for RDF Query Processing

Lei Gai, Wei Chen and Tengjiao Wang
*School of Electrical Engineering and Computer Science
Peking University. Beijing, P.R.China
Email: {lei.gai, pekingchenwei, tjwang}@pku.edu.cn

arXiv:1510.07749v1 [cs.DB] 27 Oct 2015

***Abstract*—RDF query optimization is a challenging problem. Although considerable factors and their impacts on query efficiency have been investigated, this problem still needs further investigation. We identify that decomposing query into a series of light-weight operations is also effective in boosting query processing. Considering the linked nature of RDF data, the correlations among operations should be carefully handled. In this paper, we present SGDQ, a novel framework that features a partition-based *Summary Graph Driven Query* for efficient query processing. Basically, SGDQ partitions data and models partitions as a summary graph. A query is decomposed into subqueries that can be answered without inter-partition processing. The final results are derived by perform summary graph matching and join the results generated by all matched subqueries. In essence, SGDQ combines the merits of graph match processing and relational join-based query implementation. It intentionally avoids maintain huge intermediate results by organizing sub-query processing in a summary graph driven fashion. Our extensive evaluations show that SGDQ is an effective framework for efficient RDF query processing. Its query performance consistently outperforms the representative state-of-the-art systems.**

I. Introduction

The Resource Description Framework(RDF) have been pervasively adopted in many fields, such as knowledge bases(DBpedia [1], YAGO [2], Freebase [3]), semantic social networks (FOAF [4], Open Graph [5]) and bioinformatics(Bio2RDF [6], UniProt [7]), etc. The available RDF data keep increasing in both size and quality, this challenges for a framework for efficient query processing. For schema-flexible RDF data, query is naturally a graph pattern matching process. Generating results in a refined data space can greatly boost the RDF query performance. To this end, a considerable exploratory methods have been proposed in the recent decade.

One class of solutions utilize the graph nature of RDF data and implement query as subgraph matching [8]–[10]. Subgraph matching only needs to explore the neighborhood data, and generates results in an incremental fashion. This avoids maintaining large intermediate results [9]. Nonetheless, the scale of RDF graph have a profound impact on query performance. The time complexity of subgraph matching increase exponentially as the graph size increase. Furthermore, as graph exploration is based on random data access. the efficiency of exploration-intensive subgraph match is guaranteed by holding the whole graph topology in memory, which is infeasible for large-scale graph. Recent researches [11], [12] introduced graph summarization in RDF query, The condense summary graph are solely used in filtering out irrelevant data in query.

The other class, which are the mainstream solutions, are based on 40+ years of relational database researches [13], [14]. In principle, RDF data is managed as relational tables, and joins are excessively utilized in query implementation. The scan and join on large table are prone to introduce large intermediate results, which can easily overwhelm the system memory and cause memory thrashing that greatly degrade query performance. Existing query optimization methods mainly focus on triple indexing techniques(e.g. [15], [16]) and query optimization methods(e.g. [17], [18]) that utilizing the typical characteristics of RDF. Using relational solutions, we argue that splitting RDF data into smaller parts that can be processed independently is also effective in optimizing query processing. On the flip side, there is no one-size-fit-all partition method, some results may span across partitions. The existence of inter-partition processing complicated the query implementation in this manner. Considering that the extensive connections among partitions can be naturally modeled using graph structure, this envisions us to introduce a summary graph and use subgraph match to drive the operations on partitions.

In this paper, we present a new RDF processing framework named SGDQ(***S**ummary **G**raph **D**irected **Q**uery*). SGDQ partitions data and decomposes a query into sub-queries, such that each sub-query can be answered independently in partitions. All sub-query processing are driven by a summary graph matching process, and the final results is generated as joins between subqueries in a match. The rational behind SGDQ is two folds. First, SGDQ alleviates the need of maintaining massive intermediate results by decomposing query process into a series of light-weight sub-queries. Second, subgraph match is effective in representing the inter-partition query processing, and it is can be efficiently implemented on a summarized graph of moderate scale.

To achieve this, two aspects in SGDQ need to be carefully considered. The first aspect is the design of physical data layout that support decomposed sub-query processing. We distinguish triples that do not contribute to graph partition and manage them in a specific manner. To facilitate summary match, we adopt a graph partition based strategy for data partition. Each partition is managed individually and the interconnection among partitions is modeled as a summary graph. The second aspect is the efficient processing of sub-queries. We propose a process that use auxiliary patterns to filter out irrelevant data in sub-query processing, and adopt a mechanism to prevent duplications in final results.

We summarize our major contributions in this paper as follows:

- We present SGDQ(***S**ummary **G**raph **D**irected **Q**uery*), a novel and effective framework for RDF query processing. SGDQ decomposes a query into a series of light-weight sub-queries, and leverages summary graph matching paradigm to drive the processing of sub-queries. SGDQ can greatly reduce the amount of data that need to maintain during query, and boost query performance by effective summary graph matching and efficient sub-query processing.
- We propose a specialized physical data layout that support SGDQ, and provide efficient operations based on this physical data layout. These operations are the building blocks of SGDQ.
- We perform extensive evaluations using two types of benchmarks. The evaluation results show that the query performance of SGDQ consistently outperforms three existing representative RDF systems.

The rest of the paper is organized as follows. *Section* II describes the principle concepts. *Section* III gives an overview of SGDQ framework. *Section* IV presents the physical data layout and implementation of sub-query processing. *Section* V details the paradigm of summary graph driven processing. The most related works is reviewed in *Section* VI , and *Section* VII presents the evaluation results. We conclude in *Section* VIII.

II. PRELIMINARIES

RDF is a W3C recommended standard for representing and publishing linked web resources. Due to its simple and schema-free nature, RDF facilitates the representation and integration of data from different domains. In RDF, A real world *resource* is identified by an Uniform Resource Identifier(URI) string. Its properties and linkages to other resources are described by a set of triples. Each triple is uniformly represented as three elements $\langle$*Subject, Predicate, Object*$\rangle$, or $\langle$s,p,o$\rangle$ for brevity. We call *p* a *relation* if *o* is a resource URI, or an *attribute* if *o* is a literal(string, number, or date, etc)[1]. A collection of such triples forms an *RDF dataset* $\mathcal{D} = \bigcup_{i=1}^{|\mathcal{D}|} t_i$. $\mathcal{D}$ can be naturally modeled as a directed, edge-labeled *RDF graph* $\mathcal{G} = \langle V, E\rangle$. Each triple $t_i \in \mathcal{D}$ corresponds to a directed labeled edge $e \in E$, which connects vertex *s* to vertex *o* with an edge value *p*.

Like SQL in relational database, users commonly resorts to declarative query for RDF data access. Adopting query language such as SPARQL [19], a conjunctive RDF query Q can be expressed as: **SELECT** $?v_1, \ldots$ **FROM** TP, where TP is a combination of triples, each one of which has a least one of *s*, *o* or *p* replaced by variables $?v$. We refers to such triples as *triple patterns*, and denotes a query as $Q = \{tp_1, \ldots, tp_k\}$. Two triple patterns are *joinable* if they implicitly share a common elements. In this paper, we assume that all predicates of *triple pattern*s may not be variables. This assumption is reasonable for that *Predicate* join is infrequent in practice, as shown in a previous study [20]. This leads to 3 types of joins, named according to the position of common elements: *s-s join*, *o-o join*, *s-o join*. By considering each tp_i as an edge and follows same way of mapping $\mathcal{D}$ to $\mathcal{G}$, Q can also be viewed as a *query graph* G^Q. Accordingly, a query can be implemented as the process of match G^Q on $\mathcal{G}$.

The notations used in this paper are listed in *Table* I. Next, we further introduce some principle concepts.

TABLE I: Notations used in this paper.

Notation	Description
$\mathcal{D}$	RDF dataset, $\mathcal{D} = \bigcup_{i=1}^{|\mathcal{D}|} t_i$, t_i in N-Triple (s,p,o).
s, *p* and *o*	denotes the value of *Subject*,*Predicate* and *Object* in a triple.
$\mathcal{D}_A, \mathcal{D}_R$	The set of *A-Triples* and *R-Triples*.
$\mathcal{M}_A, \mathcal{M}_R$	The global dictionary of $\mathcal{D}_A$ and $\mathcal{D}_R$.
$\mathcal{G}$	The RDF graph of $\mathcal{D}$, $\mathcal{G} = \langle V, E\rangle$.
Q	Query, as a set of k triple patterns. $Q = \{tp_1, \ldots, tp_k\}$
$\mathcal{G}_R$	The Resource-Linkage graph of $\mathcal{D}$, $\mathcal{G}_R = (V_R, E_R)$.
G_R^Q	The Resource-Linkage graph of Q, $G_R^Q = \langle V_R^Q, E_R^Q\rangle$.
$\mathcal{P}$	A partition set of $\mathcal{D}$. $\mathcal{P} = \{\mathcal{P}_1, \ldots, \mathcal{P}_n\}$. $\bigcup_{i=1}^{n} \mathcal{P}_i = \mathcal{D}_R$.
$\mathcal{SG}$	The summary data graph that models $\mathcal{P}$. $\mathcal{SG} = \langle V_S, E_S, f_v\rangle$.
Q^p	A decomposition of G_R^Q. $Q^p = \{Q_1^p, \ldots, Q_m^p\}$. $Q_i^p = (V_i^p, E_i^p)$, $\bigcup_{i=1}^{m} E_i^p = E_R^Q$
$\mathcal{TG}^Q$	The summary query graph. $\mathcal{TG}^Q = \langle V_T^Q, E_T^Q, f_v^q\rangle$.
$\mathbb{P}_{1\text{-}UHC}$	The graph partition strategy .
Op_{sp}	The processing a sub-query on a partition.
Op_{ipj}	The inter-partition join processing .

A. Triple Classification

Considering that the *relation* and *attribute Predicate* show different characteristics, $\mathcal{D}$ and Q fall into two categories define as follow:

*Definition 1 (**R/A-Triple,R/A-Pattern**):* A triple $t_i \in \mathcal{D}$ is called a *R-Triple*(an *A-Triple*) if its *p* is a *relation*(an *attribute*). The set of all *R-Triple*s(or *A-Triple*s) in $\mathcal{D}$ is denoted as $\mathcal{D}_R$(or $\mathcal{D}_A$). Analogously, A triple pattern $tp_i \in Q$ is called a *R-Pattern*(an *A-Pattern*) if its *p* is a *relation*(an *attribute*), and the set of *R-Pattern*(*A-Pattern*) in Q is denoted as Q_R(or Q_A).

Intuitively, a $tp_i \in Q_R (\in Q_A)$ can only be match to $t_j \in \mathcal{D}_R (\in \mathcal{D}_A)$. Using *Definition*1, we can get a compact graph model defines as follow:

*Definition 2 (**Resource-Linkage(RL) Graph**):* The graph representation of $\mathcal{D}_R$(or Q_R) is named as the Resource-Linkage graph of $\mathcal{D}$(or Q), or *RL graph* for brevity, denoted as $\mathcal{G}_R$ (or G_R^Q).

According to *Definition*2, $\mathcal{G}_R$ (or G_R^Q) only models the linkage information of resources. That provides a compact representation of $\mathcal{G}$ (or G^Q). Intuitively, *Property* 1 holds.

*Property 1 (**RL Graph Match**):* Following the *subgraph match* defined in [21], if there exists a subgraph g in $\mathcal{G}$ that match G^Q, consequently, the *RL graph* of g is a match of G_R^Q.

Example 1: All triples of an RDF dataset $\mathcal{D}$ are listed in *Fig.*1(a), $\mathcal{D}_A$ in the left , and $\mathcal{D}_R$ in the right. $\mathcal{D}$ represents part of a twitter-like semantic social network, which has 3 social accounts, 2 posted contents and 4 replies. *Fig.*1(b) shows the equivalent *RDF graph* $\mathcal{G}$, and *Fig.*1(c) is the RL-graph of $\mathcal{G}$. Given a query Q with 12 *triple patterns*(shown in *Fig.*2(a)), its *query graph* G^Q is represented as *Fig.*2(b). For $\mathcal{D}$, there are two matched triple sets of Q, lists in *Fig.*2(c). Therefore, the result of Q are {*u1,u3,p1,r3,r4*}, and {*u1,u3,p1,r2,r1*}.

[1]There are some exceptions. Following the semantics of RDF schema, predicate *rdf:type* is widely used to link instances with classes. Thus in this paper, *rdf:type* is classified as an *attribute*. Furthermore, some predicates can also be *attribute*. The selection of these *attribute*-predicate is dataset specific.

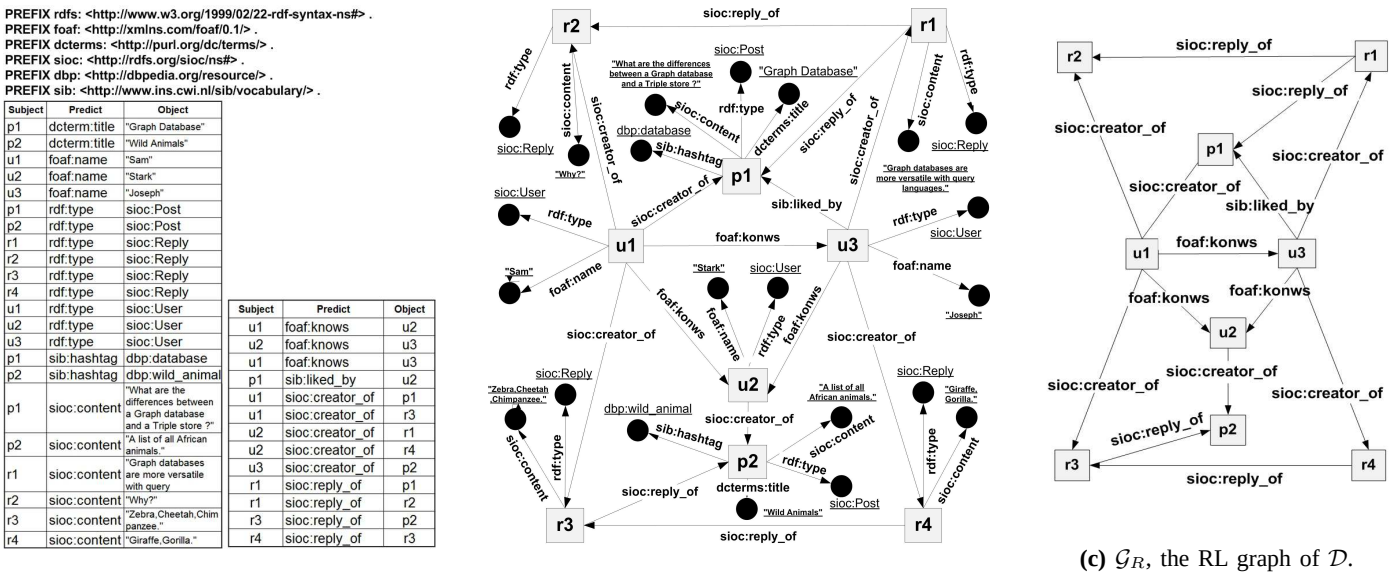

(a) An examplar RDF dataset, $\mathcal{D}$.

(b) $\mathcal{G}$, the RDF graph of $\mathcal{D}$.

(c) $\mathcal{G}_R$, the RL graph of $\mathcal{D}$.

Fig. 1: An exemplar RDF dataset $\mathcal{D}$, its data graph $\mathcal{G}$ and RL-graph $\mathcal{G}_R$

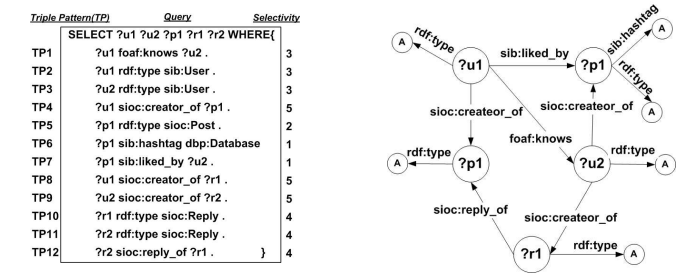
Triple Pattern(TP) Query Selectivity
SELECT ?u1 ?u2 ?p1 ?r1 ?r2 WHERE{
TP1 ?u1 foaf:knows ?u2 . 3
TP2 ?u1 rdf:type sib:User . 3
TP3 ?u2 rdf:type sib:User . 3
TP4 ?u1 sioc:creator_of ?p1 . 5
TP5 ?p1 rdf:type sioc:Post . 2
TP6 ?p1 sib:hashtag dbp:Database 1
TP7 ?p1 sib:liked_by ?u2 . 1
TP8 ?u1 sioc:creator_of ?r1 . 5
TP9 ?u2 sioc:creator_of ?r2 . 5
TP10 ?r1 rdf:type sioc:Reply . 4
TP11 ?r2 rdf:type sioc:Reply . 4
TP12 ?r2 sioc:reply_of ?r1 . } 4

(a) Q.

(b) G^Q.

Matched triple set (I)

{ (u1 foaf:knows u3), (u1 rdf:type sib:User), (u3 rdf:type sib:User),(u1 sioc:creator_of p1), (p1 rdf:type sioc:Post),(p1 sib:hashtag dbp:DataBaseSE), (p1 sib:liked_by U2), (u1 sioc:creator_of r3), (u3 sioc:creator_of r4),(r3 rdf:type sioc:Reply), (r4 rdf:type sioc:Reply), (r4 sioc:reply_of r3) }

Matched triple set (II)

{ (u1 foaf:knows u3), (u1 rdf:type sib:User), (u3 rdf:type sib:User), (u1 sioc:creator_of p1), (p1 rdf:type sioc:Post), (p1 sib:hashtag dbp:DataBaseSE), (p1 sib:liked_by u2), (u1 sioc:creator_of r2), (u3 sioc:creator_of r1), (r1 rdf:type sioc:Reply), (r2 rdf:type sioc:Reply), (r1 sioc:reply_of r2) }

(c) Lists of matched triples of Q in $\mathcal{D}$.

Fig. 2: An exemplar query and matched triples.

B. Summary Graph and Query

For an RDF graph $\mathcal{G}$, graph partition allows triples that are close to each other in $\mathcal{G}$ to be managed in the same partition. Thus, a large $\mathcal{G}$ can be summarized into a small graph that keeps the schematic topological information among partitions. We introduce a kind of summary graph defines as follow:

*Definition 3 (**Summary RDF Graph**):* Given an RDF dataset $\mathcal{D}$, a set $\mathcal{P} = \{\mathcal{P}_1, \ldots, \mathcal{P}_n\}$ is the partition set of its corresponding RDF graph $\mathcal{G} = \langle V, E\rangle$. A summary RDF graph $\mathcal{SG} = \langle V_S, E_S, f_v\rangle$ is a labeled, undirected multi-graph, where $V_S = \{\mathcal{P}_1, \ldots, \mathcal{P}_n\}$, and E_S is consist of: **i)** edges that connects two vertexes in V_S if there exists a connecting edge between two partitions; **ii)**loop edges for each vertexes in V_S, $e = (v, v)$,$v \in V_s$. f_v maps all distinct *p* in $\mathcal{P}_i$ as vertex label.

Query decomposition is extensively used in converting a complex query into a set of simple sub-queries. A query Q_R can be decomposed into a set of sub-queries, and each sub-query is a subgraph of the original query graph G^Q. We use the term *subgraph* and *sub-query* interchangeably in this paper. Based on query decomposition, we introduce the definition of *transformed query graph* as follow:

*Definition 4 (**Transformed Query Graph**):* Given a query Q and its corresponding query graph $G^Q = \langle V^Q, E^Q\rangle$. Let one of its decomposition as $Q^p = \{Q^p_1, \ldots, Q^p_m\}$, where $Q^p_i = \langle V^p_i, E^p_i\rangle$, $\bigcup_{i=1}^{m} E^p_i = E^Q$. A *transformed query graph* $\mathcal{TG}^Q = \langle V^Q_T, E^Q_T, f^q_v\rangle$ for Q is a vertex-labeled, undirected graph, where $V^Q_T = \{Q^p_1, \ldots, Q^p_m\}$, and edge $(Q^p_i, Q^p_j) \in E^Q_T$ if their exists common vertices between Q^p_i and Q^p_j. f^q_v maps all distinct *p* in Q^p_i as vertex label.

Example 2: We partition the $\mathcal{G}_R$ in *Fig.* 1(c) into two partitions, $\mathcal{P}^o = \{\mathcal{P}^o_1, \mathcal{P}^o_2\}$, show as *Fig.* 4(a). Using $\mathcal{P}^o$, we can model a *Summary RDF Graph* $\mathcal{SG}$ as *Fig.* 4(c). For the query in *Fig.* 2(a), its query *RL-Graph* Q_R can be decomposed into four subgraphs $Q^p = \{Q^p_1, \ldots, Q^p_4\}$, show as *Fig.* 4(b). Using Q^p, we model a *transformed query graph* $\mathcal{TG}^Q$ as *Fig.* 4(d).

C. Summary Graph Match

Lets OP_{sp} refers to the operation of a sub-query on a data partition. More precisely, the processing of a $Q^p_i \in Q^p$ on a $\mathcal{P}_a \in \mathcal{P}$ is denoted as $Q^p_i(\mathcal{P}_a)$, and its result set is denoted as $R(Q^p_i(\mathcal{P}_a))$. To deal with the connection between sub-queries, we introduce the concept of intra-partition join, expressed as OP_{ipj}. For two OP_{sp} namely $Q^p_i(\mathcal{P}_a)$ and $Q^p_j(\mathcal{P}_a)$, $1 \le i, j \le m$, $1 \le a, b \le n$, a OP_{ipj} is the join processing between $R(Q^p_i(\mathcal{P}_a))$ and $R(Q^p_j(\mathcal{P}_b))$, denoted as $OP_{ipj}(Q^p_i(\mathcal{P}_a), Q^p_j(\mathcal{P}_b))$, and the result set is denoted as $R(Q^p_i(\mathcal{P}_a), Q^p_j(\mathcal{P}_b))$. We define the summary graph match of $\mathcal{TG}^Q$ on $\mathcal{SG}$ as follow:

*Definition 5 (**Summary Graph Match**):* Given a $\mathcal{TG}^Q = \langle V^Q_T, E^Q_T, f^q_v\rangle$ and a $\mathcal{SG} = \langle V_S, E_S, f_v\rangle$, where $V^Q_S =$

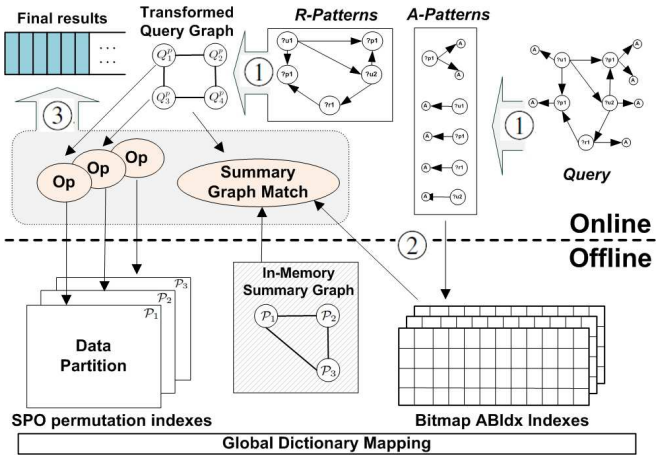

Fig. 3: The Architecutre of SGDQ framework.

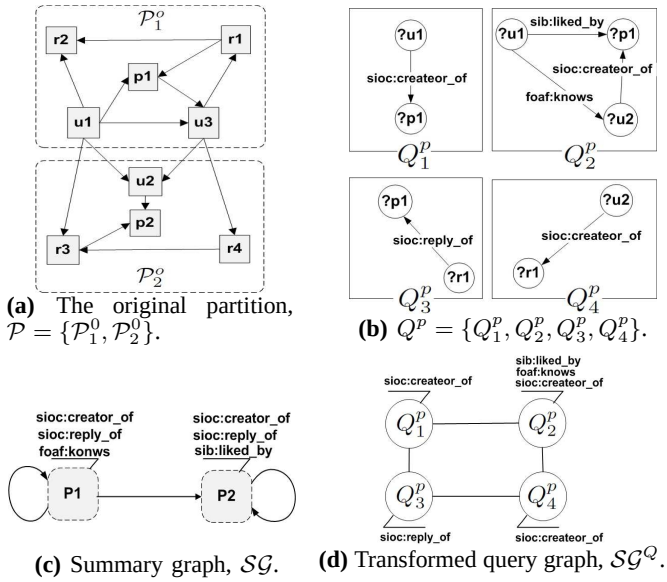

(a) The original partition, $\mathcal{P} = \{\mathcal{P}_1^0, \mathcal{P}_2^0\}$.

(b) $Q^p = \{Q_1^p, Q_2^p, Q_3^p, Q_4^p\}$.

(c) Summary graph, $\mathcal{SG}$.

(d) Transformed query graph, $\mathcal{SG}^Q$.

Fig. 4: An example of data partition and query rewrite.

$\{Q_1^p, \ldots, Q_m^p\}$,$V_S = \{\mathcal{P}_1, \ldots, \mathcal{P}_n\}$, a *summary graph match* is an injective function $M : V_T^Q \rightarrow V_S$, such that: **i) vertex match:** $\forall Q_i^p \in V_T^Q$, $f_v^q(Q_i^p) \subseteq f_v(M(Q_i^p))$, and $R(Q_i^p(\ M(Q_i^p))) \neq \varnothing$, and **ii) edge injection:** $\forall (Q_i^p, Q_j^p) \in E_T^Q$, $(M(Q_i^p), M(Q_j^p)) \in E_S$, and $R(Q_i^p(M(Q_i^p)), Q_j^p(M(Q_j^p))) \neq \varnothing$.

III. Overview of SGDQ

Fig. 3 gives an overview of SGDQ framework. A dataset $\mathcal{D}$ are offline indexed to facilitate online query processing.

A. Offline Processing

Given an RDF dataset $\mathcal{D}$. according to the *p* of each triple, $\mathcal{D}$ is divided into two parts, $\mathcal{D}_A$ and $\mathcal{D}_R$. $\mathcal{D}_A$ is managed on disk using a set of bitmap-based indexes(*Section* IV-A). Besides, $\mathcal{D}_R$ is modeled as a *data RL-graph* $\mathcal{G}_R$, and partitioned into a set of subgraphs $\mathcal{P}^o = \{\mathcal{P}_1^o, \ldots, \mathcal{P}_n^o\}$ using METIS algorithm [22](denoted as $\mathbb{P}_{METIS}$). We call $\mathcal{P}^o$ as an *original partitions*. $\mathbb{P}_{METIS}$ is an edge-cut graph partition method, it divides a graph into a given number(there it is n) of roughly vertices-balanced connected subgraphs $\mathcal{P}_j^o = (V_j^o, E_j^o), 1 \leq j \leq n$, with minimum number of cutting edges. Therefore, we can model $\mathcal{P}^o$ as a summary data graph $\mathcal{SG}$ according to *Definition* 3. $\mathcal{SG}$ is maintained as an in-memory adjacency list to facilitate graph exploration. Each partition in $\mathcal{P}^o$ is further expanded using a partition strategy which duplicates boundary triples, and final partitions $\mathcal{P} = \{\mathcal{P}_1, \ldots, \mathcal{P}_n\}$ are generated. Each $\mathcal{P}_i \in \mathcal{P}$ is managed as a self-governed triple store. All triples in $\mathcal{D}_R$ are encoded globally using a global dictionary. This enables inter-partition processing. (*Section* IV-B).

B. Runtime Query Processing

Once a user interactively submit a query Q, the query processing is implemented in three steps.

Step ①, *Query preprocessing*. Under the premise of no *Predicate* join in Q , Q is divided into Q_A and Q_R, judged by whether *p* of each triple pattern is a *relation* or an *attribute*. Q_R is modeled as a $\mathcal{G}_R^Q$, while Q_A are grouped by *Subject*. By decomposing $\mathcal{G}_R^Q$ into set of sub-queries $Q^p = \{Q_1^p, \ldots, Q_m^p\}$, we model a $\mathcal{TG}^Q$ (*Section* IV-C).

Step ②, *A-Pattern processing*. Foremost, each group of Q_A is processed using bitmap indexes, and the candidate set of the corresponding vertex in $\mathcal{G}_R^Q$ are generated for further used in OP_{sp}.(*Section* IV-A).

Step ③, *Summary graph driven processing*. Based on data partition and query decomposition, the query processing is decomposed as a set of light-weight OP_{sp}. Follow the paradigm of match $\mathcal{TG}^Q$ on $\mathcal{SG}$, We organize the execution of OP_{sp} and perform OP_{ipj} between two OP_{sp}. The final results is generated in an incremental fashion, and stops until no more matches could be found. (*Section* V-A).

IV. Data Organization and Operations

In this section, we explain how to decompose a costly query processing a series of light-weight Op_{sp} operations. A common practice is to introduce data partition and query reformulation. Our implementation is efficient in two aspects. On one side, we identify two types of query pattern to simplify a query. Utilize the characteristics of *A-Pattern*, we introduce bitmap index for *A-Triples* and implement efficient bitwise operation for *A-Pattern* processing(*Section* IV-A). On the other side, we introduce a data partition strategy(*Section* IV-B) and corresponding query decomposition method, and implement an efficient Op_{sp}(*Section* IV-C).

A. Bitmap Indexes for A-Triples and A-Pattern processing

Managing and querying *A-Triple*s in a specific manner brings two main benefits. First, is avoids unnecessary duplication of *A-Triple*s in data partitions. This simplifies data partition. Second, it employs efficient bitmap-based bitwise operations that take advantage of the characteristics of *s-s join* for *A-Pattern*s.

Bitmaps are commonly used in databases as they can be efficiently compressed and processed. Using bitmap in RDF data, *s-s joins* on a common *Subject* can be translate to simple bitwise AND operations. This saves a potentially large number of joins compared with canonical relational query processing. For this reason, some RDF researches [14], [16] have also considered the extensive adoption of bitmap indexes. To encode data as bitmap, we build two types of global dictionary, one for $\mathcal{D}_R$ and the other for $\mathcal{D}_A$, denoted as $\mathcal{M}_R$ and $\mathcal{M}_A$ respectively. $\mathcal{M}_R$ holds the mapping of an unique ID, namely *VertexID*, to a *s* or an *o* in $\mathcal{D}_R$(i.e. $v \in V_R$). $\mathcal{M}_R$ is built according to the *original partitions*. We consecutively allocate a *VertexID* to each vertex grouped by partitions, and store all *VertexID*s in an sorted order in $\mathcal{M}_R$. Similarly, $\mathcal{M}_A$ is the mapping of an *AttributeID* to an *o* in $\mathcal{D}_A$. All *AttributeID*s are managed as a sorted list in $\mathcal{M}_A$. The partial order lists of *VertexID* and *AttributeID* in $\mathcal{M}_R$ and $\mathcal{M}_A$ are further used in bitmap indexes construction.

Based on encoding data globally using $\mathcal{M}_R$ and $\mathcal{M}_A$, we design and implement a set of compressed column-oriented indexes that dedicated for *A-Triples* management and *A-Pattern*s processing. Consider there are two types of *A-Pattern*,

Pattern-I (?s,p,o), where *o* is a designated URI or Literal. *Pattern-II(?s,p,?o)* and *(s,p,?o)*, where *?o* is a variable.

For *Pattern-I*, we design a bitmap of size $|V_R|$, referred to as *BitVector*, to represent the existence of all *VertexID* for a given combination of *o* and *p*. In a *BitVector*, a bit at position i corresponds to the i-th *VertexID* in $\mathcal{M}_R$, and is set to 1 if triple (*VertexID,p,o*) $\in \mathcal{D}_A$, or 0 otherwise. There are totally $|\mathcal{D}_A|$ such *BitVector*s. They are grouped by *p*, and each group is sorted and indexed on *o* using B+-tree to facilitate index lookup. Consider that the size of *BitVector*s (i.e. $|V_R|$) can be quite large, a common way in practice is using compression to take advantage of the sparse nature of the vectors. Similar to [16], we adopts D-Gap compression[2]. This allows bitwise operations to be performed against compressed data, which eliminates the cost of *AttributeVector* decompression. Such index can be viewed as a special implementation of **POS** permutation index. We name it as $\mathsf{ABIdx}_{\mathrm{POS}}$.

For *Pattern-II*, we design two kinds of index, namely, $\mathsf{ABIdx}_{\mathrm{comp}}$ and $\mathsf{ABIdx}_{\mathrm{PSO}}$. In $\mathsf{ABIdx}_{\mathrm{comp}}$, each *p* maintains a bitmap which stores the bitwise AND results of all *BitVector*s of this *p* in $\mathsf{ABIdx}_{\mathrm{POS}}$. This bitmap captures all *A-Triples* that have this *p*. $\mathsf{ABIdx}_{\mathrm{PSO}}$ can be viewed as a special implementation of **PSO** index, and implemented the same way as $\mathsf{ABIdx}_{\mathrm{POS}}$, where a *BitVector* of size $|\mathcal{M}_A|$ is arranged according to the sorted list of *AttributeID* in $\mathcal{M}_A$.

At query time, for a $v \in V_R^Q$ of a given Q, let TP_A be the set of *A-Pattern*s that have v as *Subject*. We introduce an operation, named *CandidateRLVertex*, that inputs TP_A and returns the existence of candidate *VertexID*s for v as a *BitVector* of size $|V_R|$. For each pattern $tp_i \in TP_A$, *CandidateRLVertex* first check the type of tp_i. If tp_i a *Pattern-I*, it retrieves $\mathsf{ABIdx}_{\mathrm{PSO}}$ according to *p* and *o* of tp_i and get a *BitVector*. Otherwise, if tp_i a *Pattern-II*, it gets a *BitVector* in $\mathsf{ABIdx}_{\mathrm{comp}}$ corresponding to the *p* of tp_i. The final result is a *BitVector* that is derived from the bitwise AND of all retrieved *BitVector*s. This *BitVector* represents the existence of all *VertexID*s. To get the corresponding *AttributeID* for the variable *?o* in *Pattern-II*, we implement an extra operation, named *RetrieveAttribute*, at the final stage of Q processing. *RetrieveAttribute* retrieves the *BitVector* in $\mathsf{ABIdx}_{\mathrm{POS}}$ according to the *VertexID* in previous results and *p* in *Pattern-II*, then retrieves $\mathcal{M}_A$ to get the corresponding *AttributeID*, and maps to its real value.

Example 3: For the *original partition* (*Fig.*4(a)) of $\mathcal{D}$ (*Fig.*1(a)), we arrange an ordered vertex list as {u1,u3,p1,r1,r2} for $\mathcal{P}_1^o$ and {u2,p2,r3,r4} for $\mathcal{P}_1^o$. In $\mathcal{M}_R$, we consecutively allocate *VertexID* 1 to 5 for the former list, and 6 to 9 for the later one. Given a query with 2 patterns, *Q*={*?s rdf:type sioc:Reply*, *?s sioc:content ?o*}, we first convert all strings into ID using $\mathcal{M}_A$ and $\mathcal{M}_R$, then retrieve $\mathsf{ABIdx}_{\mathrm{PSO}}$ according to *Predicate rdf:type* and *Object sioc:Reply*, and get a *BitVector* of [000110011], Meanwhile,$\mathsf{ABIdx}_{\mathrm{comp}}$ is retrieved using a key of *Predicate sioc:content*, and [001110111] is got. By performing bitwise AND between two retrieved *BitVector*, we get a result *BitVector* [000110011]. To get the final results, we get all 1-bit corresponding *VertexID*s using $\mathcal{M}_R$, and retrieve $\mathsf{ABIdx}_{\mathrm{POS}}$ for each *VertexID* and *Predicate sioc:content*. and get a result *BitVector* each time. The final result can be derived by checking the 1-bit corresponding *AttributeID* in $\mathcal{M}_A$.

B. R-Triples Partition and Index

Next, we introduce how $\mathcal{D}_R$ is partitioned. In order to explain what type of query Q can be implemented without inter-partition operations, we introduce the concept of *query coverage*. Given a set of partitions $\mathcal{P} = \{\mathcal{P}_1, \ldots, \mathcal{P}_n\}$ generated by a partition strategy $\mathbb{P}$ over $\mathcal{D}$. If a query Q can be completely answered in $\mathcal{P}$ without inter-partition operations, we call that partition $\mathcal{P}$ can *cover* Q .If $\mathbb{P}$ can completely *cover* all query graphs in which the shortest path between two arbitrary vertices $< \phi$, we call ϕ as the *query coverage* of $\mathbb{P}$. To increase ϕ, a common way is to introduce triple replication in each $\mathcal{P}_i$, $1 \leq i \leq n$. We defines *replication factor* $\alpha = \frac{|\mathcal{D}_A| + \sum_{i=1}^{n} |\mathcal{P}_i|}{|\mathcal{D}|}$ to quantitatively measure such replication. There is a tradeoff between ϕ and α . For example, $\mathbb{P}_{METIS}$ have $\phi = 0$, $\alpha = 1$. This means although $\mathbb{P}_{METIS}$ does not introduces any replicated triples, it must deal with inter-partition query processing for an arbitrary query. Increase ϕ can results in an exponential increase of replicated triples, which led to an extra storage overhead.

To achieve a reasonable ϕ with acceptable α, in our work, we utilize a $\mathbb{P}$ named *Undirected One Hop Cover*(1-UHC) partition strategy, denoted as $\mathbb{P}_{1\text{-}UHC}$. $\mathbb{P}_{1\text{-}UHC}$ follows the idea of *undirected 1-hop guarantee* introduced by [23]. To elaborate, based an *original partition* $\mathcal{P}_j^o = (V_j^o, E_j^o) \in \mathcal{P}^o$ for $\mathcal{G}_R$, let $V_c = \{v | (v, v_i) \in E_j^o \vee (v_i, v) \in E_j^o, v_i \in V_j^o, v \notin V_j^o\}$ denotes the set of cutting edges in $\mathcal{P}_j^o$. $\mathbb{P}_{1\text{-}UHC}$ generates expanded partitions $\mathcal{P}_j = (V_j, E_j)$ following these steps:

1) Adds all cutting edges $\forall e \in \{e = (v_a, v_b) \in \mathcal{G}_R | (v_a \in V_j^o \wedge v_b \in V_c) \vee (v_a \in V_c \wedge v_b \in V_j^o)\}$ to E_j;
2) Adds all edges that connect cutting edges $\forall e \in \{e = (v_a, v_b) \in \mathcal{G}_R | (v_a \in V_c \wedge v_b \in V_c)\}$ to E_j;
3) Assigns $\forall v \in V_c$ to V_j.

$\mathbb{P}_{1\text{-}UHC}$ differs from [23] in that it solely applied on $\mathcal{G}_R$. As $\mathcal{G}_R$ is more compact than $\mathcal{G}$, this leds to a smaller α.

[2]http://bmagic.sourceforge.net/dGap.html

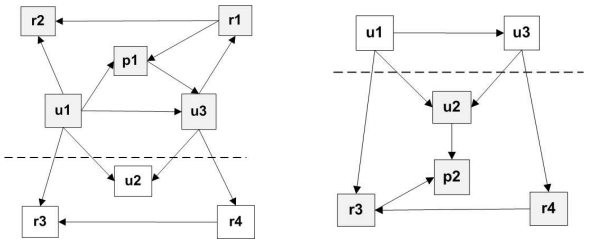

(a) $\mathcal{P}_1$. **(b)** $\mathcal{P}_2$.

Fig. 5: A raw partition and its 1-OHC partitions.

Example 4: For $\mathcal{G}_R$ in *Fig.*1(c), an *original partition* $\mathcal{P}^o = \{\mathcal{P}_1^o, \mathcal{P}_2^o\}$ generated by $\mathbb{P}_{METIS}$ are shown as *Fig.*5(a). After apply $\mathbb{P}_{1\text{-}UHC}$, we get partition $\mathcal{P} = \{\mathcal{P}_1, \mathcal{P}_2\}$ shown as *Fig.*5.

For management of triples in a partition, we independently index and organize them as a set of permutation indexes. This is supported by most state-of-the-art triple store. There we adopted RDF-3X [13]. To enable joins between partitions, for each partition, we use the *VertexID* in $\mathcal{M}_R$ to globally encode the *s* and *o* of all triples. We additionally maintain two metadata of a partition. One represents all *VertexID* in a $\mathbb{P}_{1\text{-}UHC}$ partition. It is stored as a *BitVector*, named as *pVector*. *pVector* encodes the existence of $v_i \in V_R$ in a partition follows the partial order of *VertexID* in $\mathcal{M}_R$. It is useful in filtering triples using *A-Pattern*s. The other represents all *VertexID*s in an *original partition*. As they are allocated consecutively in $\mathcal{M}_R$, we only need to store the start and end *VertexID*. This is useful in avoiding results duplication. We discuss the use of indexes and metadata in the subsequent *Section* IV-C.

Selection of *n* is an empirical problem. A premise of our work is that each partitions can be efficiently processed without too much memory thrashing. This led to a rough estimation that the size of a partition and the possible intermediate results for arbitrary query processing can be a well fit to the restricted memory. In our evaluation, we adopt $n = 500$.

C. Query decomposition and Sub-query Processing

Next, we explain how to decompose G_R^Q into a set of subgraphs $Q^p = \{Q_1^p, \ldots, Q_m^p\}$ according to $\mathbb{P}_{1\text{-}UHC}$, and the light-weight implementation of *OP*$_{sp}$ considering *A-Pattern*s. Furthermore, we introduce a mechanism in *OP*$_{sp}$ to avoid generating duplications in final query results.

According to *Definition* 4, a *transformed query graph* $\mathcal{TG}^Q$ is constructed by transform each Q_i^p into a $\mathcal{TG}^Q$ vertex, and establish an edge between two $\mathcal{TG}^Q$ vertices if their corresponding subgraphs in Q^p share a common vertex in Q_R^G. As each $Q_i^p \in Q^p$ can by covered by $\mathbb{P}_{1\text{-}UHC}$, and $\mathbb{P}_{1\text{-}UHC}$ has $\phi = 2$, there are only two types of sub-queries :

Type-I *simplex triple pattern*: a single pattern in G_R^Q.
Type-II *complete graph pattern*: a set of triple patterns that forms an undirected complete graph in G_R^Q. The most simple case is a triangle of three patterns.

The generation of a transformed query graph can be implemented in a greedy fashion. We start with selection of a vertex v with the maximum degree in $G_R^Q = \langle V_R^Q, E_R^Q \rangle$, then traverse through all v's neighbour vertices to check whether a complete graph exists. If exists, we add a vertex to $\mathcal{TG}^Q$, and add an edge between existing vertices in $\mathcal{TG}^Q$ to newly added vertex if they share a common vertex. If all complete graph in G_R^Q that contains v has been detected, we remove current vertex v from V_R^Q and all edges of all complete graph from E_R^Q. We iterate the above process until there are no edges in E_R^Q. Obviously in worst case, $\mathcal{TG}^Q$ can have $|E_R^Q|$ vertices.

As the basic building block of SGDQ, an *OP*$_{sp}$, $Q_i^p(\mathcal{P}_j)$, is the sub-query $Q_i^p \in Q^p, 1 \leq i \leq m$ on partition $\mathcal{P}_j \in \mathcal{P}, 1 \leq j \leq n$. Optimization of *OP*$_{sp}$ can greatly enhance the overall query performance. Furthermore, as $\mathbb{P}_{1\text{-}UHC}$ introduce overlapping among partitions, in *OP*$_{sp}$ a mechanism is needed to tackle the problem of latent duplications in final results.

Before explain the implementation of *OP*$_{sp}$, we first introduce a process, named as *filter*. The *filter* uses *A-Pattern*s to prune the matched triple set of a *R-Pattern* in a $\mathcal{P}_j$. Given a *R-Pattern*, let v_o be its *Subject* vertex, v_s be its *Object* vertex. Each vertex may have *A-Pattern*s connect to it, and the candidate set of v_s and v_o , denoted as $CR(v_s)$ and $CR(v_o)$, can be got using *CandidateRLVertex* operation. For all triples in $\mathcal{P}_j$ that match a *R-Pattern*, *filter* eliminates triples that theirs *Subject*(or *Object*) are not contained in $CR(v_o)$(or $CR(v_s)$). The set containment check is also based on *BitVector*s, and can be efficiently implemented in $O(1)$ time.

For *Type-I* sub-query, which is a match of simplex triple pattern, *OP*$_{sp}$ is implemented in a *scan-filter* fashion. The *scan* is a scan of a pattern-specific permutation index. In *filter*, $CR(v_o)$ is derived by a bit-wise AND between the *CandidateRLVertex* result *BitVector* of v_o and the partition's *pVector*. For $CR(v_s)$, we distinguish between two cases. If current *OP*$_{sp}$ corresponds to the last matched vertex in the match *order*(discuss in detail in *Section* V-A), denoted as the *last OP*$_{sp}$ for brevity, $CR(v_s)$ is derived by a range selection of the *CandidateRLVertex* result *BitVector* of v_s. If *OP*$_{sp}$ is not the *last OP*$_{sp}$, $CR(v_s)$ is derived the same way as $CR(v_o)$. Both range data and *pVector* is stored as metadata in a partition. By only keep the triples that have *Subject* in the overlapping-free *original partition*, duplications in the final results can be avoided.

For *Type-II* sub-query, we implement *OP*$_{sp}$ in a *scan-filter-join-verify* fashion. In essence, as each $\mathcal{P}_j$ is managed by a standalone RDF-3X triple store, we can utilize the query mechanism provided by RDF-3X for *OP*$_{sp}$. More precisely, RDF-3X uses a *scan-join* mechanism for query processing. A query is decomposed into a set of triple patterns. The candidate set of each triple pattern is generated by a scan of a corresponding permutation index. All candidate sets are joined following an optimized order and get the final query results. We further add a *filter* operation before join really happens. At *filter* stage, $CR(v_s)$(and $CR(v_o)$) is derived by bit-wise AND between the *CandidateRLVertex* result *BitVector* of v_s(and v_o) and the partition's *pVector*. The *filter* introduces *A-Pattern*s to filter irrelevant triple from participate in join operation, this design optimize the original *scan-join* processing. To avoid final results duplication, we introduce a *verify* stage right after *join* process to prune the results of *scan-filter-join* processing if current *OP*$_{sp}$ is the *last OP*$_{sp}$. The basic idea behind *verify* is that it keeps the sub-query results that tend to be generated by this *OP*$_{sp}$. This tendency is measured by that the result

graph in $\mathcal{G}_R$ have more vertices in the *original partition* generated by $\mathbb{P}_{METIS}$. It is rational on that $\mathbb{P}_{METIS}$ follows an iterative *coarsen-partition* processing scheme, complete graph $K_n, n \geq 4$ are more likely to be treated as whole in each *coarsen* stage. This means that $K_n, n \geq 4$ are always allocated in a partition. To deal with exceptions, we assume that at most one vertex in a K_n is involved in partition. To illustrate, for K_3, which is a triangle, we keep the results of OP_{sp} that have two vertices in its *original partition*. And for K_4, we keep the results that have three vertices in its *original partition*. This is implemented as an range check of all *VertexID* of a result using the range information in the metadata of this partition.

As each *Type-II* sub-query needs to determine its own query plan, the total time of query plan selection can sum up to a high cost. In RDF query processing, cardinality estimation is difficult for the lack of schema. A heuristic-based query plan needs less time to find a sub-optimal plan. We adopt a heuristic-based query plan selection method for each OP_{sp} follows [24]. Considering the size of data partition, a sub-optimal plan is still acceptable for an OP_{sp}.

V. Summary Graph Driven Query(SGDQ) Processing

In this section, we explain how SGDQ use summary graph match process to direct OP_{sp} processing.

A. Summary Graph Match

According to *definition* 5, the summary graph match is a special case of subgraph isomorphism problem. This problem can be solved using a generic backtracking method [21]. *Algorithm* 1 outlines an implementation of backtracking-based summary graph match. It determines a match order of vertices in $\mathcal{TG}^Q$(*Line* 3), then recursively calls a match procedure to find matched $\mathcal{TG}^Q$ in $\mathcal{SG}$. The match process directs the processing of Op_{sp}(*Line* 6-20), and outputs query results of Q on $\mathcal{G}$. As SGDQ follows the general backtracking paradigm, we only explain the specific features used in SGDQ.

GetCandidateRLVertex(*Line* 2) process all *A-Pattern*s in Q grouped by common *Subject*. For each query *RL-graph* vertex $v \in V_R^Q$, *GetCandidateRLVertex* gets the the candidate vertices in G_R as a *BitVector*(v). If v is a variable with no connected *A-Pattern*s, all bits in *BitVector*(v) are assigned to 1. If v is an URI of a *VertexID*, the corresponding bit in *BitVector*(v) are assigned to 1, and 0 elsewhere. Otherwise, a *CandidateRLVertex* operation is called to process all *A-Pattern*s that connect to v. □

DetermineMatchOrder(*Line* 3). In principle, the *order* is determined as an order of DFS graph traversal in $\mathcal{TG}^Q$. Thus the selection of the starting vertex in $\mathcal{TG}^Q$ is important. We adopt heuristics that considers the output of *GetCandidateRLVertex* and the the statistics of $\mathcal{D}$. These heuristics are listed as follows, and can be used in combination or separately.

Heuristic 1. Choose the order according to the selectivity of all $Q_i^p \in V_T^Q, 1 \leq i$. The selectivity of Q_i^p is roughly estimated as the smallest candidate set size of $v \in Q^p$ that has the smallest candidate set in V_R, or the smallest selectivity of *Predicate*s of all *R-Pattern*s in Q_i^p. The former can be derived using *GetCandidateRLVertex* function, and the candidate set of Q_i^p in $\mathcal{SG}$ is determined by a range check of all 1-bits in *BitVector*(v). The later is based on the statistics of $\mathcal{D}$, which are pre-computed as metadata. The $Q_i^p \in V_T^Q$ with smallest selectivity is chosen as the starting vertex. After starting vertex is selected, in the DFS traversal that follows, high selective vertex goes first.

Heuristic 2. Given $Q_i^p \in V_T^Q$ that have been arranged in order, if $Q_j^p \in V_T^Q$ is a neighbour of Q_i^p in $\mathcal{TG}^Q$, and merge-join can be applied between them, then arrange Q_j^p as the next in order.

Heuristic 3. If Heuristic 1 and 2 can not apply, select Type-II Q_i^p into order with priority. □

Algorithm 1: SGDQ processing.

Input: A query Q, $\mathcal{D}$, with $\mathcal{D}_A$ in $\mathsf{ABIdx}_{\mathrm{POS}}$, $\mathsf{ABIdx}_{\mathrm{comp}}$ and $\mathsf{ABIdx}_{\mathrm{PSO}}$ indexes, and $\mathcal{D}_R$ as $\mathcal{P} = \{\mathcal{P}_1, \ldots, \mathcal{P}_n\}$, summary RDF graph $\mathcal{SG}$ of $\mathcal{D}$.
Output: All results of Q on $\mathcal{D}$.
begin
1 Construct $\mathcal{TG}^Q = \langle V_G^Q, E_G^Q \rangle$, $V_G^Q = \{Q_1^p, \ldots, Q_m^p\}$;
2 $CR(v)$ =GetCandidateRLVertex($\forall v \in V_G^Q$ $\forall$*A-Patterns*$\in Q$);
3 $order$ = ***DetermineMatchOrder***($\mathcal{TG}^Q$,$\mathcal{SG}$);
4 ***SubgraphMatch***($\mathcal{TG}^Q$,$\mathcal{SG}$,M,$order$,1);
5 RetrieveAttribute(); //special steps for *Pattern-II* processing;
6 SubgraphMatch($\mathcal{TG}^Q$,$\mathcal{SG}$,M,$order$,d) **if** $d == |V_T^Q|$ **then**
7 report R; //R stores intermediate join results;
8 v_s = GetNextCandidate($order$);
9 $C(v_s, M(v_{prev}))$ = ***ExploreCandidates***($\mathcal{TG}^Q$,$\mathcal{SG}$,v_s,$CR(v_s)$,$order$);
10 **foreach** $v \in C(v_s, M(v_{prev}))$, v *is not matched* **do**
11 $R_v \leftarrow v_s(v)$; //Op_{sp} of v_s on partition v.;
12 **if** $R_v = \varnothing$ **then**
13 $CR(v_s) = CR(v_s) - v$; continue ;
14 R= $Op_{ipj}(\mathbb{S}(d-1), R_v)$;
15 **if** $R = \varnothing$ **then**
16 continue ;
17 $\mathbb{S} \leftarrow R$; ForwardTrack($\mathbb{S}$,v_s,v);
18 ***SubgraphMatch***($\mathcal{TG}^Q$,$\mathcal{SG}$,M,$order$,$d+1$);
19 $\mathbb{S} = \mathbb{S} - R$; BackTrack($\mathbb{S}$,$v_s$,$v$);
20 **return**;

SubgraphMatch(*Line* 6-20) is a recursive function. It starts as gets a current query vertex v from *order*(*Line* 8), and obtains a set of candidate vertices of $M(v)$ in $\mathcal{SG}$, denoted as $C(v, M(v_{prev}))$, which is the set of neighbour vertices[3] of the previous matched vertices $M(v_{prev})$. The vertex match in $\mathcal{SG}$ is judged by vertex label match, and corresponding bit of $M(v) \in C(v, M(v_{prev}))$ in *BitVector*($M(v)$) is 1-bit(*Line* 9).

SubgraphMatch iterates over all candidate vertices $M(v) \in C(v, M(v_{prev}))$ (*Line* 10-19). In each iteration, *SubgraphMatch* calls an Op_{sp}, i.e. $v(M(v))$, and gets its results $R(v(M(v)))$(*Line* 11). If $R(v(M(v))) = \varnothing$, v is removed by setting the corresponding bit in *BitVector*(v) to 0(*Line* 13). This avoids unnecessary repeated execution of Op_{sp} that do not satisfy vertex math in *Definition* 5. For each edge between v and all v_{prev} in $\mathcal{TG}^Q$, an $Op_{ipj}(v(M(v), v_{prev}(M(v_{prev})))$ is called to check for edge injection(*Line* 14). As the search order follows a $\mathcal{TG}^Q$ graph traversal order, in most cases Op_{ipj} are implemented using hash join. We use order-preserving merge-joins for Op_{ipj} if possible, considering *Heuristic 2* in *DetermineMatchOrder*. If the result set

[3]Considering the existence loop edges for each vertex $\mathcal{SG}$, the neighbours contains the vertex itself.

$R(v(M(v), v_{prev}(M(v_{prev})))$ is not $\varnothing$, the join results are cached in a structure $\mathbb{S}$(*Line* 17). $\mathbb{S}$ enables backtracking to previous matched vertex, and the cached intermediate results in $\mathbb{S}$ can be reused for all vertices in the candidate region $C(v, M(v_{prev}))$.

Next, *SubgraphMatch* recursively calls another *SubgraphMatch* with an incremental in depth d(*Line* 18). If d equals to the number of vertex in $\mathcal{TG}^Q$, this means a match of $\mathcal{TG}^Q$ is generated, and currently cached R can be output as part of final results(*Line* 7). If Op_{ipj} returns $\varnothing$, *SubgraphMatch* iterate to the next candidate if exists, otherwise backtrack to the next candidate vertex in $\mathcal{SG}$ for $v_{prev} \in \mathcal{TG}^Q$ in order. □

*Property 2 (**Query Completeness and Non-Redundancy**):* Let R_{SGDQ} and R be the result set of generated by Q on $\mathcal{D}$ and SGDQ respectively. If $r \in R$, then $r \in R_{SGDQ}$ holds.

This property can be proofed by contradiction. We omit the details in this paper.

B. Cost Analysis

Given $\mathcal{D}_A$, and $\mathcal{D}_R$ as $\mathcal{P}$ of n partitions. For a query Q, assume that in Q there are totally t_I *Pattern-I A-Pattern*s, t_{II} *Pattern-II A-Pattern*s, and Q is decomposed into m sub-queries, of which m_I *Type-I*, and m_{II} *type-II*. The average cost of *Type-I* and *Type-II* sub-query are $Cost(Op_{sp})$ and $Cost(Op_{sp})$ respectively.

We first give the cost analysis of *A-Patterns*. As there are total $|\mathcal{D}_A|$ *BitVector*s, the cost of *CandidateRLVertex* operation is the time complexity of index lookup using a composite key of *p* and *o*, which is $O(log|\mathcal{D}_A|)$. Similarly, the cost of *RetrieveAttribute* operation is $O(log|V_R|)$. Consequently the aggregated cost of all *A-Pattern*s processing sum up to $O(t_I \cdot (log|\mathcal{D}_A|) + t_{II} \cdot log|V_R|)$

The *R-Pattern*s processing follows summary graph matching paradigm. Conventional optimization methods of subgraph matching are focused on establishing graph topological indexes to effectively prune candidate vertices. Such optimizations can not work for *Algorithm* 1, as we still needs to actually take joins to determine whether two partitions are joinable. Thus exhaustive exploration of all feasible neighbours is adopted in *Algorithm* 1. This is acceptable as $\mathcal{SG}$ size is relative small. To this end, In worst case SGDQ needs $\sum_{i=1}^{m-1} |V_S|^i$ times of Op_{sp}, and $|E_S| \cdot \sum_{i=1}^{m-1} |V_S|^i$ times of Op_{ipj}. The cost of *R-Pattern*s processing can sum up to $(\sum_{i=1}^{m-1} |V_S|^i) \cdot Cost(Op_{sp}) + (|E_S| \cdot \sum_{i=1}^{m-1} |V_S|^i) \cdot Cost(Op_{ipj})$.

SGDQ only needs to maintain intermediate result at each recursive call, which is generated on m partitions. This significantly reduces the cost to keep vast amount of immediate results generated by scans and joins on large dataset.

VI. Related Works

RDF query optimization has attracted interests from both industrial and research domain, and considerable methods, either relational or graph-based, centralized or distributed, have been proposed. In this section, we only list the most related to our approach and compare them with ours.

Graph-based approaches. RDF query in native graph storage is implemented as subgraph matching. Graph pattern matching has been extensively studied [21], [25]. Following these, [9], [10] use fine-grained graph exploration for RDF query. This need to maintain the entire graph in memory. [8], [11] introduce structural summaries of graph as disk-based structural index. Matched results are generated using a filter-and-refinement strategy. SGDQ constructs a summary graph based on graph partitions, which is most similar to [12]. SGDQ differs from [12] in that the summary graph is used in SGDQ for arrange sub-query processings, while it is used in [12] for pruning unnecessary data before join processing.

Relational approaches. A naïve way to manage RDF data is adopting relational database to store all triples in a single 3 column table. Query is implemented in a *scan-join* fashion [12], [13], [26]. Optimizations of naïve approach are mainly from three aspects. **i)Using permutation index to make join and scan themselves efficient.** [15] first introduce the concept of permutation index. Hexastore [26] exhaustively indexes all SPO permutations and stores them directly in sorted B+-tree. RDF-3X [13] further introduces extra indexes of all SPO attributes projections. Such permutation indexes enables efficient range scan and merge-join at the cost of data redundancy. SGDQ extensively use permutation indexes in manage each partitions. **ii)Reducing the number of joins needed.** Property table [27] clustering relevant *predicate*s together in a table schema. This helps in eliminating *s-s join*s, but requires a previous knowledge about query. BitMat [16] represents property tables as 3 dimension bitmaps, and replace joins with bitwise operations which share common elements at a dimension. SGDQ is different in that it classifies triples and adopts bitmaps only for *A-Triple*s. This eliminates all *s-s join*s related to *A-Pattern*s, and avoids complex considerations of other types of joins when using bitmap index. **iii)Filtering out irrelevant inputs of join**. RDF-3X adopted a technique named SIP(Sideway Information Passing) to dynamically determine the unnecessary data blocks that can be skipped during scan [28]. RG-index [29] introduces a relational operator *filter* that using frequent graph pattern to prune triples that can not match. SGDQ uses filter in each Op_{sp} and shares the same objective as [28], [29], but it uses the results of *A-Pattern*s processing in filtering.

Data partition are widely used in shared-nothing distribute RDF systems [23] [12]. They follow the principles of evenly distribute data to each computing node so that each partitions can be processed in parallel to boost the overall query performance. The number of partitions are generally equal to the number of computing nodes, and the goal of optimization is to design effective partition strategy to minimize the cost of distribute inter-partition joins. To this end, the most common adopted partition strategies are edge-based hash [30] and graph partition [23]. SGDQ adopts partition in a totally different way. It partition data in order to decompose original query into a series of light-weight processing, the number of partitions depends on the dataset size.

VII. Evaluation

In this section, we empirically evaluate the RDF query processing performance on two kinds of large-scale benchmark data. The goals of our evaluations are **i)** We show the correctness and the competitive performance of SGDQ on basic benchmark queries (contain at most 6 triple patterns) over

the state-of-the-art *triple stores*(*Section* VII-C); **ii)** We show the superior performance of SGDQ on complex queries that involve up to 16 triple patterns, and analyze the effect of our methods(*Section* VII-D, VII-E).

A. Environment, Competitors and Datasets

All experiments were performed on a server with Debian 7.4 in 64-bit Linux kernel, two Intel Xeon E5-2640 2.0GHz processors and 64GB RAM. The query response time was measured in milliseconds. SGDQ is developed in C++, and is compiled using GCC-4.7.1 with -O3 optimization flag. We used METIS 5.1[4] to get the *original partition* of RDF graph, and RDF-3X 0.3.8[5] for partition management and support of sub-query processing. For fair comparison, We also used the dictionary to convert between the string value of *s*,*o*,*p* and the identifiers used in processing, and presented results in its original string value form.

Competitors. There exists a wide choice on the state-of-the-art *triple store*. We chosen three publicly available and representative ones. We list these competitors as:

RDF-3X maintains all possible permutation of *s*,*p* and *o* as indexes, and use B+-tree to facilitate index lookup. Query of a single pattern is implemented as an index scan operation, and merge join are extensively leveraged to join between two patterns. Together with index-specific query optimization techniques [13], [28], RDF-3X remains a competitive state-of-the-art *triple stores*.

TripleBit[6] organizes data in a compact bit matrix, and vertically partition the matrix into sub-matrices following predicates dimension and sorted subject/object dimension. Such compact data representation and partition can minimizes the size of intermediate results, thus theoretically results in a query performance promotion [14].

Virtuoso 7[7] is a commercial-of-the-shelf column store that support RDF management. It exploit existing relational RDBMS techniques and add functionality to deal with features specific to RDF data. It also adopted bitmap index to take advantage of the property that many triple share the same predicate and object [31].

Dataset. We adopted two different kinds of benchmarks.

LUBM[8] is a widely used benchmark in both academical researches and industrial applications. It provides a university database where the components, such as university, departments, professors, students etc, can be polynomially generated. The size of dataset scales according to the number of universities. We generated two dataset, they are identified as LUBM-8000 and LUBM-20480 according to the number of universities. As all components in the data are generated in a proportional fashion, the RDF graph of a LUBM dataset can be regarded as a random graph.

SNIB[9] provides the semantic dataset of a twitter-like social network, which includes resources like users, post, reply, tags and comments, etc. The dataset scales according to the number of users. We generate a SNIB dataset of 15000 users using S3G2 [32], identified as SNIB-15000. As it models a social network, the RDF graph of SNIB-15000 is a power-law graph.

Remind that to support the original LUBM queries, the standard way is to inference on RDF Schema *subclasses* and *subproperties*. We used the inference engine in Virtuoso 7 in our experiment, and the inferred triples were also managed by each triple store. The statistics of dataset are listed in *Table* II.

TABLE II: Statistics of data collection .

Dataset	# Triples, $\lvert\mathcal{D}\rvert$		RL-Graph, $\mathcal{G}_R$		#P
	# Totals	# Inferred	$\lvert V\rvert$	$\lvert E\rvert$	
LUBM-8000	1,970,947,330	869,030,729	173,806,051	529,614,087	18
LUBM-20480	4,844,909,074	2,024,718,666	444,851,202	1,355,446,182	18
SNIB-15000	1,503,766,205	27,618,900	125,779,052	875,721,431	48

Queries. For LUBM dataset, we used the 14 queries provided by LUBM benchmark. Furthermore, we manually constructs 6 complex queries on LUBM and and 7 queries on SNIB dataset. These complex queries contains more triple patterns than the original ones. All Queries are listed in the *Appendix*.

Consider that the cache mechanism of operating system have a profound impact in query execution, we measure the performance using both cold cache mode and warm cache mode. In cold cache mode, we purged the OS file system cache each time before running a query[10]. Each query was executed in cold cache mode and warm cache mode for 10 times in a successive manner. The result in each mode was reported as the arithmetic mean and was round to millisecond[11].

B. Effects of Partitions

We adopted $n = 500$, i.e. the RL-graph of each dataset was partitioned into 500 subgraphs using METIS. The effects of $\mathbb{P}_{METIS}$ and $\mathbb{P}_{1\text{-}UHC}$ is given in *Table* IV, manifested as the number of cutting edges and *replication factor* α respectively. By managing *A-Triples* individually, $\mathbb{P}_{1\text{-}UHC}$ have far smaller duplicated triples compared with that of the *un-two* strategy show in [23].

We list the physical storage characteristics in *Table* V[12]. Basically, RDF-3X costs the most space, as it needs to maintain all 6 SPO permutation indexes and 9 aggregated indexes of binary and unary projection. Even though SGDQ use RDF-3X to manage all partitions of *R-Triples*, but the bitmap index of *A-Triple*s contributes to a more compact storage. Consider the proportion of *A-Triples* in a dataset(as shown in *Table* II), SGDQ shows a more efficient storage compared with RDF-3X.

C. LUBM benchmark queries

We report the response time of query on LUBM-8000 and LUBM-20480 dataset in *Table* IIIa and *Table* IIIb respectively. Generally speaking, SGDQ achieve a comparable query

[4]http://glaros.dtc.umn.edu/gkhome/fetch/sw/metis/metis-5.1.0.tar.gz
[5]https://rdf3x.googlecode.com/files/rdf3x-0.3.8.zip
[6]http://grid.hust.edu.cn/triplebit/TripleBit.tar.gz
[7]https://github.com/openlink/virtuoso-opensource
[8]http://swat.cse.lehigh.edu/projects/lubm/
[9]http://www.w3.org/wiki/Social_Network_Intelligence_BenchMark

[10]In Debian OS, we purged its file system cache using commands: **/bin/sync; echo 3 > /proc/sys/vm/drop_caches**
[11]Actually, in warm cache mode, a query was executed 11 times consecutively. The first execution was for cache warm-up, and its result was removed from averaging.
[12]Virtuoso 7 manage all dataset in a single database instance. The storage characteristic of Virtuoso 7 is measured by the increased database file size after each dataset is loaded.

TABLE III: Benchmark query response time on LUBM dataset(in milliseconds).
★: Wrong in query result size. **N/A**: Query does not end in 12 hours.

(a) LUBM-8000 dataset.

	L1	L2	L3	L4	L5	L6	L7	L8	L9	L10	L11	L12	L13	L14
Cold Cache.														
RDF-3X	377	1,044,390	372	456	377	2,508,477	541	1,040	517,775	385	497	440	24,479	1,828,117
Virtuoso7	**7**	25,193	3	**7**	**5**	2,450,552	**4**	628	225,441	2	5	55	4,381	1,981,602
TripleBit	115	91,085★	**2**	62	51	3,891,150	37	109★	423,428	**0.3**	**3**	**5**	626	3,122,737
SGDQ	92	**11,790**	42	56	49	**924,535**	102	**325**	**94,702**	22	63	92	**271**	**592,374**
Warm Cache.														
RDF-3X	2	910,820	2	5	7	2,196,851	7	439	109,162	2	186	19	6,174	1,553,394
Virtuoso7	2	24,033	1	3	4	2,381,088	2	427	223,551	1	2	21	1,592	1,568,623
TripleBit	**0.2**	39,337★	**0.3**	3	3	3,823,121	2	102★	178,381	**0.2**	2	**4**	582	3,045,244
SGDQ	2	**8,015**	1	**2**	**3**	**892,722**	**2**	**221**	**76,118**	1	**2**	32	**64**	**573,189**
#Results	4	2528	6	34	719	83,557,706	67	7,790	2,178,420	4	224	15	37,144	63,400,587

(b) LUBM-20480 dataset.

	L1	L2	L3	L4	L5	L6	L7	L8	L9	L10	L11	L12	L13	L14
Cold Cache.														
RDF-3X	1,327	2,383,700	377	533	394	6,265,536	511	1,028	843,811	387	1,845	497	72,766	4,310,605
Virtuoso7	**7**	41,169	**5**	**8**	**5**	6,661,334	**59**	888	425,757	9	8	19	4,291	5,587,519
TripleBit	124	187,712★	98	112	89	8,224,312	137	182★	739,002	**1**	**3**	**6**	1,127	7,811,369
SGDQ	76	**24,327**	64	75	68	**1,224,376**	112	**122**	**172,411**	21	72	107	**433**	**1,502,331**
Warm Cache.														
RDF-3X	2	2,330,893	3	5	7	5,781,561	7	401	830,944	2	362	14	38,988	3,813,701
Virtuoso7	2	42,521	1	**3**	4	6,253,872	4	402	440,644	3	3	24	3,759	4,532,172
TripleBit	**0.2**	56,005★	**0.3**	2	**3**	8,102,716	5	169★	527,032	**0.5**	3	**4**	639	7,692,721
SGDQ	2	**12,073**	2	5	4	**1,123,549**	**4**	**23**	**122,395**	1	**2**	17	**117**	**1,422,801**
#Results	4	2,528	6	34	719	213,817,916	67	7,790	2,703,043	4	224	15	95,522	162,211,567

TABLE IV: Partition characteristics

Dataset	# cutting edges	α
LUBM-8000	23,624,351	1.23
LUBM-20480	61,518,672	1.21
SNIB-15000	58,823,356	1.52

TABLE V: Physical storage characteristics (in GB).

Dataset	Raw N3 file	RDF-3X	Virtuoso7	TripleBit	SGDQ
LUBM-8000	343	124	58	57	92
LUBM-20480	848	283	129	147	184
SNIB-15000	203	97	42	49	84

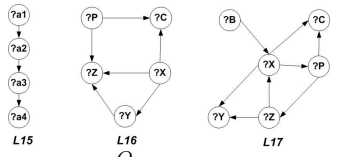

Fig. 6: The G_R^Q of *L15*, *L16* and *L17*

performance in comparison to its competitors using standard benchmark queries.

Query *L1*, *L3*, *L4*, *L5*, *L7*, *L8*, *L10*, *L11* exhibit the same characteristics. They all have a high-selective *A-Pattern*, and their result sets are quite small and are independent of the dataset size. Among them, *L4* is a bit more complex, it contains 3 *Pattern-II A-Pattern*s. SGDQ needs an extra bitmap lookup to retrieve the variable of these *A-Pattern*. Due to their simple nature, they can not benefit from SGDQ optimizations. In fact, the response time of all these queries for both SGDQ and all competitors are within 1 second for both cold cache and warm cache. It is hard for a user to feel the difference. Although Query *L13* also have a high-selective *A-Pattern*, the result set of *L13* increase as the dataset size increase. For cold cache query, RDF-3X needs to scan the large indexes and cache the data at each query time. SGDQ outperforms RDF-3X by 4.1(*L1* on LUBM-8000) to 168(*L13* on LUBM-20480) times. This proves the effectiveness of SGDQ in avoiding the cost of scan large indexes by manage and query data in partitions. As for warm cache query, the candidate triples of high selective patterns can be held in cache. As SGDQ introduce extra cost of bitmap search and partition selection, SGDQ has no advantages over competitors.

For the time consuming queries, i.e. *L2*, *L6*, *L9* and *L14*, SGDQ outperforms all competitors by 1.7(*L2* for Virtuoso 7 on LUBM-20480) to 98(*L2* for RDF-3X on LUBM-20480) times for cold cache, and by 2.3(*L9* for TripleBit on LUBM-8000) to 193(*L2* for RDF-3X on LUBM-20480) times for warm cache. For *L6* and *L14*, SGDQ is efficient in that they only need to look up the bitmap indexes of *A-Triples*, the cost is due to the dictionary operation for string and ID conversation. For *L2* and *L9*, theirs *RL-graph* is a triangle. This means that they can be answered without inter-partition processing, and theirs $\mathcal{TG}^Q$ have only one vertex. Intuitively, the sum of all OP_{sp} are far better than execute the query as a whole.

D. Complex queries on LUBM

As standard LUBM benchmark queries can not fully evaluate the performance of SGDQ, we further design 3 more complex queries on LUBM dataset. These queries show a more complex structural information. The G_R^Q of *L15-17* are shown as *Figure* 6. We make additional 3 queries by adding a highly selective *A-Pattern* to each of them, and get $L15_{\text{sel}}$-$L17_{\text{sel}}$ accordingly. These queries are listed in *Appendix A*. Intuitively, the $\mathcal{TG}^Q$ of *L15-17* have 3,4,4 vertices and 2,3,4 edges individually.

We report the response time of query on LUBM-8000 and LUBM-20480 dataset in *Table* VI, in both cold cache and warm cache mode. In fact, for complex query on large-scale dataset, this evaluation shows that only Virtuoso 7, which is a commercial-off-the-shelf system, is qualified. RDF-3X fails

TABLE VI: Complex queries response time on LUBM-20480 dataset(in milliseconds).
★: Wrong in query result size. **N/A**: Query does not end in 12 hours.

	LUBM-8000						LUBM-20480					
	L15	L16	L17	$L15_{sel}$	$L16_{sel}$	$L17_{sel}$	L15	L16	L17	$L15_{sel}$	$L16_{sel}$	$L17_{sel}$
Cold Cache.												
RDF-3X	666,179	989,280	N/A	3,783	14,968	83,241	1,642,307	N/A	N/A	30,810	75,919	113,143
Virtuoso7	528,767	25,533	**12,492**	**81**	731	474	1,364,557	41300	**24,285**	**82**	954	544
TripleBit	1,942,030★	189,630★	95,959★	3,815★	13,379★	25,145	1,188,727★	535,515★	1,643,845★	5,900★	19,433★	9,731
SGDQ	**249,532**	**21,727**	16,392	122	**322**	**215**	**852,380**	**39,271**	28,183	158	**334**	**292**
Warm Cache.												
RDF-3X	483,763	966,009	N/A	1,499	6,519	47,877	1,184,803	N/A	N/A	3,783	14,968	92,775
Virtuoso7	522,486	24,428	12,072	63	383	371	1,345,170	41,013	23,626	**83**	951	444
TripleBit	1,800,790★	171,661★	59,726★	3,812★	13,227★	19,281	1,037,115★	456,658★	1,582,875★	4,748★	5,890★	6,708
SGDQ	**239,953**	**21,116**	**11,771**	**54**	**92**	**114**	**833,402**	**32,632**	**22,690**	117	**172**	**149**
#Results	6,717,143	1,116	21	720	4	3	17,200,845	1,116	21	720	10	3
$\#OP_{sp}$	579,364	76,943	1,543	27	22	7	591,271	77,233	1,454	27	21	7

TABLE VII: Complex queries response time on SNIB-15000 dataset(in milliseconds).

	S1	S2	S3	S4	S5	S6	S7
RDF-3X	2,396,685	16,227	84,584	8,792	N/A	5,214,652	449,401
Virtuoso7	4,134,085	19,178	4,650	5,643	23,386	829,220	805,995
SGDQ	**1,953,027**	**12,454**	**1,021**	**2,352**	**3,922**	**227,932**	**66,407**
$\#OP_{sp}$	1,762,736	5,931	69	1,293	1,127	18,556	500

to return for *L17* on LUBM-8000 and *L16*, *L17* on LUBM-20480. TripleBit returns with wrong number of results for all queries except $L17_{\text{sel}}$. For non-selective queries *L15-17*, SGDQ consistently outperform others in both warm cache and cold cache mode.

To further investigate this problem, we record the total counts of OP_{sp} operations for each query processing. We can see that *L15* generates more OP_{sp}, as each sub-queries of *L15* is a simplex pattern pattern. SGDQ implements such OP_{sp} and OP_{ipj} at the magnitude of approximately 1 ms, this contributes to the overall query performance promotion. Compare the result of *L15* on LUBM-8000 and LUBM-20480, we can see that as the data size increase, the average cost of OP_{sp} increase accordingly. This is due to expansion of partition size. Although *L17* generates far more less OP_{sp}, but most of its OP_{sp} are triangles, and the average execution time is at a magnitude of 10 to 20 ms. This evaluation shows the efficiency of OP_{sp} in LUBM dataset.

E. Complex queries on SNIB

We also evaluated the query performance on a SNIB benchmark dataset, which is generated as a social network of 15000 users. SNIB dataset presents a more connected graph structure compared with LUBM dataset, thus we can introduce more complex queries to evaluate the effect of SGDQ. We define 7 SNIB queries as listed in *Appendix B*. Their G_R^Q are shown as *Figure* 6. We report the response time of query on SNIB-15000 dataset in *Table* VII. All queries were executed in warm cache mode. As TripleBit is inclined to to return wrong number of results, we do not select it in this evaluation. In essence, SGDQ consistently outperforms all competitors as the queries are more complex.

To elaborate, query *S1* and *S4* shows a similar characteristic. All the OP_{sp} are *Type-I* sub-query. This cause the number of OP_{sp} grows exponentially with the number of edges in G_R^Q. SGDQ shows a superior performance due to the efficiency of OP_{sp}. Query *S3* and *S7* have a triangle G_R^Q, and their $\mathcal{TG}^Q$ contains only one vertex. They differ in that *S7* does not have

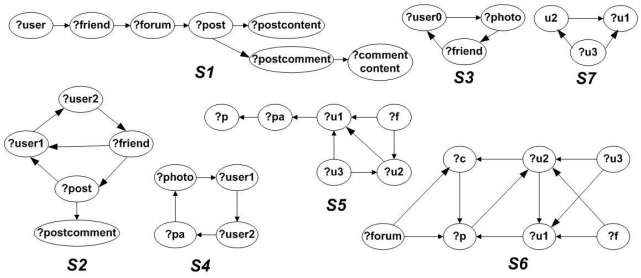

Fig. 7: The G_R^Q of 7 SNIB queries

a selective *A-Pattern* as *S3*. *S7* needs to explore all partitions. Like in LUBM evaluation, SGDQ also performs well on such queries. Query *S2*, *S5* and *S6* feature the processing of complex OP_{ipj} between *Type-II* sub-queries. RDF-3X performs the worst due to excessive number of joins, even fails for *S5*. The superior performance of SGDQ on these queries shows that all OP_{ipj} are also efficiently processed in SGDQ. These complex queries benefit from the simplification of query processing using query decomposition and data partition.

VIII. CONCLUSION

We have presented SGDQ, a ***S**ummary **G**raph **D**riven **Q**uery* framework for RDF query processing. SGDQ is based on efficient execution of sub-queries on partitioned data. It utilizes the effectiveness of subgraph matching in representing the inter-partition query processing, and avoids the problem of maintaining huge intermediate results in canonical relational query methods. To enable SGDQ, we introduced specialized physical data layout and implemented efficient operations in SGDQ. We have performed an extensive evaluation on the query performance of SGDQ compared with three competitive RDF systems. We also noticed that in current work, the exhaustive exploration of neighbourhood in summary graph match introduces many useless inter-partition joins. In the future, we plan to investigate the mechanism to prune unnecessary partitions by maintaining fine-grained inter-partition data in

the summary graph.

REFERENCES

[1] "DBpedia," http://wiki.dbpedia.org/.

[2] "YAGO," www.mpi-inf.mpg.de/yago/.

[3] "Freebase," https://www.freebase.com/.

[4] "The Friend of a Friend (FOAF) project," http://www.foaf-project.org/.

[5] "Fackbook Open Graph," http://ogp.me/.

[6] "Bio2RDF, linked data for the life sciences," http://http://bio2rdf.org/.

[7] "UniProt," ftp://ftp.uniprot.org/pub/databases/uniprot/current_release/.

[8] L. Zou, J. Mo, L. Chen, M. T. Özsu, and D. Zhao, "gstore: Answering sparql queries via subgraph matching," *PVLDB*, vol. 4, no. 8, pp. 482–493, 2011.

[9] K. Zeng, J. Yang, H. Wang, B. Shao, and Z. Wang, "A distributed graph engine for web scale rdf data," in *PVLDB*, 2013, pp. 265–276.

[10] J. Kim, H. Shin, W.-S. Han, S. Hong, and H. Chafi, "Taming subgraph isomorphism for rdf query processing," *PVLDB*, vol. 8, no. 11, pp. 1238–1249, 2015.

[11] T. Tran, G. Ladwig, and S. Rudolph, "Managing structured and semistructured RDF data using structure indexes," *IEEE Trans. Knowl. Data Eng.*, vol. 25, no. 9, pp. 2076–2089, 2013.

[12] S. Gurajada, S. Seufert, I. Miliaraki, and M. Theobald, "TriAD: a distributed shared-nothing RDF engine based on asynchronous message passing," in *SIGMOD*, 2014, pp. 289–300.

[13] T. Neumann and G. Weikum, "The RDF-3X engine for scalable management of RDF data," in *PVLDB*, vol. 19, no. 1, 2010, pp. 91–113.

[14] P. Yuan, P. Liu, B. Wu, H. Jin, W. Zhang, and L. Liu, "Triplebit: A fast and compact system for large scale rdf data," *PVLDB*, vol. 6, no. 7, pp. 517–528, 2013.

[15] A. Harth and S. Decker, "Optimized index structures for querying rdf from the web," in *LA-WEB*, 2005, pp. 10 pp.–.

[16] M. Atre, V. Chaoji, M. J. Zaki, and J. A. Hendler, "Matrix "bit" loaded: A scalable lightweight join query processor for rdf data," in *WWW*, 2010, pp. 41–50.

[17] M. Stocker, A. Seaborne, A. Bernstein, C. Kiefer, and D. Reynolds, "Sparql basic graph pattern optimization using selectivity estimation," in *WWWW*, 2008, pp. 595–604.

[18] A. Gubichev and T. Neumann, "Exploiting the query structure for efficient join ordering in SPARQL queries," in *EDBT*, 2014, pp. 439–450.

[19] "SPARQL," http://http://www.w3.org/TR/rdf-sparql-query/.

[20] M. Arias, J. D. Fernndez, and M. A. Martnez-Prieto, "An empirical study of real-world sparql queries," in *USEWOD@WWW*, 2011.

[21] J. Lee, W.-S. Han, R. Kasperovics, and J.-H. Lee, "An in-depth comparison of subgraph isomorphism algorithms in graph databases," in *PVLDB*, vol. 6, 2013, pp. 133–144.

[22] G. Karypis and V. Kumar, "METIS – unstructured graph partitioning and sparse matrix ordering system, version 2.0," Tech. Rep., 1995.

[23] J. Huang, D. J. Abadi, and K. Ren, "Scalable SPARQL querying of large RDF graphs," in *PVLDB*, vol. 4, no. 11, 2011, pp. 1123–1134.

[24] P. Tsialiamanis, L. Sidirourgos, I. Fundulaki, V. Christophides, and P. Boncz, "Heuristics-based query optimisation for sparql," in *EDBT*, 2012, pp. 324–335.

[25] J. Cheng, J. X. Yu, and P. S. Yu, "Graph pattern matching: A join/semijoin approach," *IEEE Trans. on Knowl. and Data Eng.*, vol. 23, no. 7, pp. 1006–1021, 2011.

[26] C. Weiss, P. Karras, and A. Bernstein, "Hexastore: sextuple indexing for semantic web data management," *PVLDB*, vol. 1, no. 1, pp. 1008–1019, 2008.

[27] K. Wilkinson, "Jena Property Table Implementation," in *SSWS*, 2006, pp. 35–46.

[28] T. Neumann and G. Weikum, "Scalable join processing on very large RDF graphs," in *SIGMOD*, 2009, pp. 627–640.

[29] K. Kim, B. Moon, and H. Kim, "R3F: RDF triple filtering method for efficient SPARQL query processing," *World Wide Web*, vol. 18, no. 2, pp. 317–357, 2015.

[30] Z. Kaoudi, K. Kyzirakos, and M. Koubarakis, "Sparql query optimization on top of dhts," in *ISWC*, 2010, pp. 418–435.

[31] O. Erling and I. Mikhailov, "Virtuoso: RDF support in a native RDBMS," in *Semantic Web Information Management - A Model-Based Perspective*, 2009, pp. 501–519.

[32] M. Pham, P. A. Boncz, and O. Erling, "S3G2: A scalable structure-correlated social graph generator," in *TPCTC*, 2012, pp. 156–172.

APPENDIX

A. LUBM Queries

PREFIX rdf: ⟨http://www.w3.org/1999/02/22-rdf-syntax-ns#⟩
PREFIX rdfs: ⟨http://www.w3.org/2000/01/rdf-schema#⟩
PREFIX ub: ⟨http://swat.cse.lehigh.edu/onto/univ-bench.owl#⟩
L1-L14: Same as LUBM Q1-Q14 used in [10] respectively.
L15:SELECT ?a1 ?a2 ?a3 ?a4 WHERE{?a1 ub:advisor ?a2. ?a2 ub:worksFor ?a3. ?a3 ub:subOrganizationOf ?a4. ?a4 rdf:type ub:University. ?a1 rdf:type ub:GraduateStudent. ?a2 rdf:type ub:FullProfessor. ?a3 rdf:type ub:Department. }
$L15_{sel}$:SELECT ?a1 ?a2 ?a3 ?a4 WHERE{?a1 ub:advisor ?a2. ?a2 ub:worksFor ?a3. ?a3 ub:subOrganizationOf ?a4. ?a1 rdf:type ub:GraduateStudent. ?a2 rdf:type ub:FullProfessor. ?a3 rdf:type ub:Department. ?a4 ub:name "University7". }
L16:SELECT ?X ?Y ?Z WHERE{?X rdf:type ub:GraduateStudent. ?Y rdf:type ub:University. ?Z rdf:type ub:Department. ?C rdf:type ub:GraduateCourse. ?X ub:takesCourse ?C. ?P rdf:type ub:FullProfessor. ?P ub:teacherOf ?C. ?P ub:worksFor ?Z. ?X ub:memberOf ?Z. ?Z ub:subOrganizationOf ?Y. ?X ub:undergraduateDegreeFrom ?Y}
$L16_{sel}$:SELECT ?X ?Y ?Z WHERE{?X rdf:type ub:GraduateStudent. ?Z rdf:type ub:Department. ?C rdf:type ub:GraduateCourse. ?X ub:takesCourse ?C. ?P rdf:type ub:FullProfessor. ?P ub:teacherOf ?C. ?P ub:worksFor ?Z. ?X ub:memberOf ?Z. ?Z ub:subOrganizationOf ?Y. ?X ub:undergraduateDegreeFrom ?Y. ?Y ub:name "University6". }
L17:SELECT ?X ?Y ?P ?B ?C WHERE{?X rdf:type ub:GraduateStudent. ?X rdf:type ub:TeachingAssistant. ?X ub:advisor ?P. ?X ub:takesCourse ?C. ?X ub:memberOf ?Z. ?X ub:undergraduateDegreeFrom ?Y. ?Y rdf:type ub:University. ?C rdf:type ub:GraduateCourse. ?Z rdf:type ub:Department. ?Z ub:subOrganizationOf ?Y. ?P rdf:type ub:AssociateProfessor. ?P ub:worksFor ?Z. ?P ub:teacherOf ?C. ?B rdf:type ub:Publication. ?B ub:publicationAuthor ?X. }
$L17_{sel}$:SELECT ?X ?Y ?P ?B ?C WHERE{?X rdf:type ub:GraduateStudent. ?X rdf:type ub:TeachingAssistant. ?X ub:advisor ?P. ?X ub:takesCourse ?C. ?X ub:memberOf ?Z. ?X ub:undergraduateDegreeFrom ?Y. ?Y rdf:type ub:University. ?Y ub:name "University786" . ?C rdf:type ub:GraduateCourse. ?Z rdf:type ub:Department. ?Z ub:subOrganizationOf ?Y. ?P rdf:type ub:AssociateProfessor. ?P ub:worksFor ?Z. ?P ub:teacherOf ?C. ?B rdf:type ub:Publication. ?B ub:publicationAuthor ?X. }

B. SNIB Queries

PREFIX foaf: ⟨http://xmlns.com/foaf/0.1/⟩
PREFIX dc: ⟨http://purl.org/dc/elements/1.1/⟩
PREFIX sioc: ⟨http://rdfs.org/sioc/ns#⟩
PREFIX sioct: ⟨http://rdfs.org/sioc/type#⟩
PREFIX sib: ⟨http://www.ins.cwi.nl/sib/vocabulary/⟩
PREFIX rdf: ⟨http://www.w3.org/1999/02/22-rdf-syntax-ns#⟩
PREFIX dbp: ⟨http://dbpedia.org/resource/⟩
S1:SELECT ?user ?commentcontent ?commentdate WHERE{?user foaf:knows ?friend. ?user rdf:type sib:User. ?friend rdf:type sib:User. ?friend sioc:moderator_of ?forum. ?post rdf:type sioc:Post. ?forum sioc:container_of ?post. ?post sioc:content ?postcontent. ?post sib:hashtag dbp:Island . ?post sioc:container_of ?postcomment. ?postcomment rdf:type sioc:Item. ?postcomment sioc:content ?commentcontent. ?postcomment dc:created ?commentdate. }
S2:SELECT ?user1 ?user2 ?friend WHERE{?user1 foaf:knows ?user2. ?user2 foaf:knows ?friend . ?friend foaf:knows ?user1 . ?user1 rdf:type sib:User . ?user2 rdf:type sib:User . ?friend rdf:type sib:User . ?friend sioc:creator_of ?post. ?post rdf:type sioc:Post . ?post sioc:content ?postcontent. ?post sib:hashtag dbp:Arrow . ?post sib:liked_by ?user1}
S3:SELECT ?user0 ?friend ?photo WHERE{?user0 rdfs:type sib:User. ?friend rdf:type sib:User. ?photo rdf:type sib:Photo . ?friend foaf:knows ?user0. ?photo sib:usertag ?friend. ?photo dbp:location "Germany". ?user0 sib:liked_by ?photo.}
S4:SELECT ?photo WHERE{?photo rdfs:type sib:Photo. ?user1 rdf:type sib:User . ?user2 rdf:type sib:User . ?pa rdfs:type sioct:ImageGallery. ?photo sib:usertag ?user1. ?photo dbp:location "Crimea". ?user1 foaf:knows ?user2. ?pa sioc:container_of ?photo. ?user2 sioc:creator_of ?pa. }
S5:SELECT ?u1 ?u2 ?p WHERE{?f sib:memb ?u1. ?f sib:memb ?u2. ?u2 foaf:knows ?u1. ?u3 foaf:knows ?u2. ?u3 foaf:knows ?u1. ?u1 sioc:creator_of ?pa. ?u1 rdf:type sib:User. ?u2 rdf:type sib:User. ?u3 rdf:type sib:User. ?pa rdf:type sioct:ImageGallery. ?pa sioc:container_of ?p. ?p rdf:type sib:Photo. ?p dbp:location "Germany".}
S6:SELECT ?u1 ?u2 WHERE{?f sib:memb ?u1. ?f sib:memb ?u2. ?u2 foaf:knows ?u1. ?u3 foaf:knows ?u2. ?u3 foaf:knows ?u1. ?u1 rdf:type sib:User. ?u2 rdf:type sib:User. ?u3 rdf:type sib:User. ?u1 sioc:creator_of ?p. ?p rdf:type sioc:Post. ?forum sioc:container_of ?p. ?forum sioc:container_of ?c. ?p sib:hashtag dbp:Island. ?c rdf:type sioc:Item. ?u2 sioc:creator_of ?c. ?c sioc:reply_of ?p. ?p sib:liked_by ?u2. }
S7:SELECT ?u1 ?u2 ?p WHERE{?u2 foaf:knows ?u1. ?u3 foaf:knows ?u2. ?u3 foaf:knows ?u1. }